

06-23-00

A

## IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Inventorship.....Erickson et al.  
 Applicant ..... Microsoft Corporation  
 Attorney's Docket No. .... MS1-586US  
 Title: Architectures For And Methods Of Providing Network-based Software Extensions

TRANSMITTAL LETTER AND CERTIFICATE OF MAILING

To: Commissioner of Patents and Trademarks,  
 Washington, D.C. 20231

From: Lance R. Sadler (Tel. 509-324-9256; Fax 509-323-8979)  
 Lee & Hayes, PLLC  
 421 W. Riverside Avenue, Suite 500  
 Spokane, WA 99201

The following enumerated items accompany this transmittal letter and are being submitted for the matter identified in the above caption.

1. Specification--title page, plus 59 pages, including 42 claims and Abstract
2. Transmittal letter including Certificate of Express Mailing
3. 16 Sheets Formal Drawings (Figs. 1-20)
4. Return Post Card

Large Entity Status ☒ [x]

Small Entity Status ☐ [ ]

Date: 6/21/00

By:

Lance R. Sadler  
 Reg. No. 38,605

CERTIFICATE OF MAILING

I hereby certify that the items listed above as enclosed are being deposited with the U.S. Postal Service as either first class mail, or Express Mail if the blank for Express Mail No. is completed below, in an envelope addressed to The Commissioner of Patents and Trademarks, Washington, D.C. 20231, on the below-indicated date. Any Express Mail No. has also been marked on the listed items.

Express Mail No. (if applicable) EL6 2435 1983

Date: 6/21/2000

By:

Lori A. Vierra

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Architectures For And Methods Of Providing Network-  
based Software Extensions**

Inventor(s):

Paul Erickson  
Oliver Fisher

ATTORNEY'S DOCKET NO. ms1-586us

## RELATED APPLICATIONS

The following patent applications are related to the present application, are assigned to the assignee of this patent application, and are expressly incorporated by reference herein:

- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Single Window Navigation Methods and Systems", bearing attorney docket number MS1-560us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Methods and Systems of Providing Information to Computer Users", bearing attorney docket number MS1-557us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Methods, Systems, Architectures and Data Structures For Delivering Software via a Network", bearing attorney docket number MS1-559us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Network-based Software Extensions", bearing attorney docket number MS1-563us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Authoring Arbitrary XML Documents Using DHTML and XSLT", bearing attorney docket number MS1-583us, and filed on the same date as this patent application;
- U.S. Patent Application Serial No. \_\_\_\_\_, entitled "Task Sensitive Methods And Systems For Displaying Command Sets", bearing attorney docket number MS1-562us, and filed on the same date as this patent application.

## TECHNICAL FIELD

This invention relates to methods and systems for providing software via a network. More specifically, the invention pertains to Internet-based delivery of software.

1     **BACKGROUND**

2             Installation of traditional PC applications requires physical media, such as a  
3     disk or CD-ROM that must be physically inserted into a computer in order for  
4     software to be loaded onto a user's computer. Typically, this process requires the  
5     user to enter settings information that can be confusing to the user. Once the  
6     software is installed, it is typically fixed in terms of its location and functionality.  
7     When the software is updated, the user must typically purchase additional physical  
8     media and repeat the installation process so that they can use the updated software.  
9     In this model, the software is fixed in its association with the computer on which it  
10    was installed. If a user moves to another computer, they will not be able to use the  
11    specific software on their machine without repeating the installation process.

12            As computing continues to evolve into the environment of computer  
13    networks such as the Internet, it has become clear that the traditional software  
14    delivery model described above is inadequate to meet the demands of consumers  
15    who desire dynamic, flexible, and adaptable software on-demand. Network-based  
16    software delivery is becoming the subject of increasing focus by those who  
17    develop and deliver software. Unlocking the potential for network-based software  
18    delivery will require smart, innovative and streamlined solutions, especially in  
19    situations where bandwidth may be limited.

20            Accordingly, this invention arose out of concerns associated with providing  
21    new software delivery models that are particularly well-suited for network-based  
22    software delivery, e.g. delivery via the Internet.

1 **SUMMARY**

2       Methods and systems for network-based software delivery are described.  
3       In one embodiment, an application program or software platform resides on a  
4       client. The program or platform is configured so that it is extensible based on  
5       software extensions that are deliverable over a network such as the Internet.  
6       Various extensions can be developed by third party developers for incorporation  
7       into the program or platform.

8       In one described embodiment, extension files that comprise a software  
9       extension are hosted on a network server such as an Internet server. These include  
10       descriptor files that describe aspects of the software extension. These descriptor  
11       files logically describe an extension to the program or platform and designate the  
12       location of other extension files. Extensions are incorporated on a client by  
13       navigating to a particular network or Internet site through which the extensions  
14       can be accessed. The files describing the extension files are downloaded on the  
15       client. These files tell the client where, when and how the particular extension can  
16       be plugged into the program or platform, as well as where to find the appropriate  
17       extension files and how to download them. The extension files are then  
18       downloaded and incorporated into the platform.

19       In one embodiment, an inventive software architecture is provided for  
20       handling and consolidating particular types of descriptive files associated with  
21       various extensions. A filtering mechanism called attachment points, is used to  
22       create handlers for the different descriptive types of files that define a software  
23       extension. Each of these handlers is known as an attachment manager.  
24       Attachment managers are provided for each extensible feature type. The  
25       attachment managers interpret data from extensions files which are supplied by

attachment points. In addition to predefined attachment managers, custom attachment managers can be created using data from attachment points. When an extension extends a particular feature type, the attachment points ensure that only the appropriate attachment manager is notified so that the feature type can be incorporated into the program or platform efficiently.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 is a high level view of a system that can be utilized in accordance with one described embodiment.

Fig. 2 is an exemplary computer system that can be utilized in accordance with the described embodiment.

Fig. 3 is a diagram of an exemplary EDF and PKG in accordance with one described embodiment.

Fig. 4 shows a portion of an EDF in accordance with the described embodiment.

Fig. 5 shows a portion of an EDF schema in accordance with the described embodiment.

Fig. 6 shows a portion of a PKG in accordance with the described embodiment.

Fig. 7 is a block diagram illustrating how file hashes can be used for versioning in accordance with one embodiment.

Fig. 8 is a block diagram that illustrates two exemplary package objects in accordance with one embodiment.

Fig. 9 is a flow diagram that describes steps in a method in accordance with one described embodiment.



## DETAILED DESCRIPTION

### Overview

The methods and systems described just below provide a mechanism by which functionality can be added dynamically to an application program or software platform. Functionalities or “extensions” as they will be referred to, can be advantageously added via a network such as the Internet. Extensions, that can implement new features or add to existing features, can be added using only a network address, such as a URL, as a basis for extension installation. That is, all of the files that comprise an extension can be maintained on the network and accessed via one or more network sites.

Extensions can be described in a variety of ways. One way utilizes a hierarchical tag-based language which facilitates handling and use of the various extensions. In one particular implementation, a software platform is provided that can incorporate various functionalities. The software platform and the inventive architecture described below enable third and fourth party developers to develop extensions for the platform that can be easily and seamlessly incorporated into the platform without having any knowledge of (or relationship with) a hosting service. A third party developer is a developer who develops an extension for the platform. A fourth party developer might be a developer who develops an extension to a third party developer’s extension. Thus, the incorporation of third and fourth party extensions is essentially a transparent process, as far as developers are concerned.

Consider for example, Fig. 1 which shows a user’s computer 100 and several so-called extension sources 102, 104, and 106. The extension sources can comprise any entity from which a software extension can be obtained via a network. In an exemplary implementation, the network can comprise the Internet,



1 although other networks (e.g. LANs and WANs) can certainly be utilized.  
2 Extension sources can include, without limitation, business entities such as retail  
3 stores that might maintain a network site. In one implementation, a user can  
4 execute software on their computer that provides an application program or  
5 software platform. In this document, the terms "application program" and  
6 "software platform" will be used interchangeably. Each of the different extension  
7 sources 102-106 can provide software extensions that can plug into the software  
8 platform that is executing on the user's machine. These extensions are deliverable  
9 via a network such as the Internet, and assist in providing applications that can be  
10 executed on the user's machine. In the described embodiment, the extensions are  
11 logically described in XML which is in line with emerging industry standards.  
12 Additionally, the use of XML assists in the future discoverability of extensions by  
13 promoting XML DOM properties on the Internet. It will be appreciated, however,  
14 that any suitable format can be used for describing the extensions, e.g. a binary  
15 description could be used.

16 Extensions can be delivered from any number of different extension  
17 sources. The inventive methods and systems provide a streamlined and organized  
18 way to handle the provided extensions. The use of XML advantageously enables  
19 efficient handling of extensions from multiple different extension sources, without  
20 unduly taxing the software components that utilize specific portions of an  
21 extension or extensions.

22 In one particular implementation, the software platform on the user's  
23 machine provides various different integrated functionalities that enable a user to  
24 accomplish different document-centric tasks. An exemplary system is described in  
25

the U.S. Patent Application entitled "Single Window Navigation Methods and Systems", incorporated by reference above.

### **Exemplary Computer Environment**

The embodiment described just below can be employed in connection with various computer systems. A computer system, for purposes of this document, can be considered as any computing device that includes some type of processor, i.e. a microprocessor, and some type of operating system. Thus, a computer system can be construed to include, without limitation, traditional desktop computers, more powerful servers, various hand-held devices such as cell phones, pocket-held computer devices and the like.

Fig. 2 shows but one exemplary computer system that can be used to implement the embodiments described herein. Computer 130 includes one or more processors or processing units 132, a system memory 134, and a bus 136 that couples various system components including the system memory 134 to processors 132. The bus 136 represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. The system memory 134 includes read only memory (ROM) 138 and random access memory (RAM) 140. A basic input/output system (BIOS) 142, containing the basic routines that help to transfer information between elements within computer 130, such as during start-up, is stored in ROM 138.

Computer 130 further includes a hard disk drive 144 for reading from and writing to a hard disk (not shown), a magnetic disk drive 146 for reading from and writing to a removable magnetic disk 148, and an optical disk drive 150 for



Computer 130 commonly operates in a networked environment using logical connections to one or more remote computers, such as a remote computer 176. The remote computer 176 may be another personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to computer 130, although only a memory storage device 178 has been illustrated in Fig. 2. The logical connections depicted in Fig. 2 include a local area network (LAN) 180 and a wide area network (WAN) 182. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets, and the Internet.

When used in a LAN networking environment, computer 130 is connected to the local network 180 through a network interface or adapter 184. When used in a WAN networking environment, computer 130 typically includes a modem 186 or other means for establishing communications over the wide area network 182, such as the Internet. The modem 186, which may be internal or external, is connected to the bus 136 via a serial port interface 156. In a networked environment, program modules depicted relative to the personal computer 130, or portions thereof, may be stored in the remote memory storage device. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Generally, the data processors of computer 130 are programmed by means of instructions stored at different times in the various computer-readable storage media of the computer. Programs and operating systems are typically distributed, for example, on floppy disks or CD-ROMs. From there, they are installed or loaded into the secondary memory of a computer. At execution, they are loaded at

1 least partially into the computer's primary electronic memory. The invention  
2 described herein includes these and other various types of computer-readable  
3 storage media when such media contain instructions or programs for implementing  
4 the steps described below in conjunction with a microprocessor or other data  
5 processor. The invention also includes the computer itself when programmed  
6 according to the methods and techniques described below.

7 For purposes of illustration, programs and other executable program  
8 components such as the operating system are illustrated herein as discrete blocks,  
9 although it is recognized that such programs and components reside at various  
10 times in different storage components of the computer, and are executed by the  
11 data processor(s) of the computer.

007290-27866560

## Extensions

An “extension”, as used in this document, will be considered to include, without limitation, software functionality and content that can be added to an application program or software platform. These additions typically provide some type of functionality that the application program may not have had before the extension was incorporated, or alter the behavior of at least one existing feature. The extension is incorporated or integrated directly into the application program in a way that changes, to some degree, the manner in which the application program behaves or operates. Extensions provide dynamically added content and can provide applications (such as an email application), plug-ins to extend existing applications (like a fax plug-in to an email application), or simply web pages, to name just a few.

In the described embodiment, extensions are described using XML, an industry-standard, text-based markup language. XML greatly facilitates the extensibility of software content. Specifically, various extensions can be authored by third parties and described in XML in a manner that enables the extensions to be readily integrated into application programs. It should be appreciated, however, the XML constitutes but one exemplary way of describing and using the extensions. Other ways can, of course, be used.

## Exemplary Extension Organization

In the described embodiment, extensions are organized in three separate but related portions: an Extension Definition File (EDF), a Package Manifest (PKG), and the code, components, or “bits” that make up or define the extension. An EDF can be, but need not be associated with a URL (Universal Resource Locator) that

1 provides a way for a client to access the EDF. By convention and choice, the PKG  
2 file is located at the same URL as the EDF. It will be appreciated that the  
3 described EDFs and PKG are each not required for the other to be used. It just so  
4 happens that, in the example that is given in this document, the two are employed  
5 together. To that end, each of these features can be separately and independently  
6 employed.

7 EDFs describe logical attachments to an application program or software  
8 platform, while PKGs specify the physical files and resources that are used in an  
9 extension. There can be a one to one correspondence between EDFs and PKGs.

10 Fig. 3 shows an exemplary organization 300 that includes an EDF 302 and  
11 a corresponding package manifest (PKG) 304. In the illustrated example, the EDF  
12 302 uses XML to describe the logical attachments or extensions to an application  
13 program. The corresponding PKG 304 specifies the physical files and resources  
14 that are associated with a particular extension. Exemplary file types are shown to  
15 the right of PKG 304 and include, without limitation, HTM, GIF, UDO, XML,  
16 DLL, and various other types of files.

### 17 18 **Extension Definition File (EDF)**

19 In the described example, an EDF is an XML file that logically describes an  
20 extension. For example, the EDF can describe HTML that makes up a user  
21 interface (UI), the objects that contain code for implementing various functions,  
22 and the like. The EDF can also contain all or part of the functionality that  
23 comprises an extension. For instance, the HTML that describes a toolbar could be  
24 incorporated directly into an EDF file, and a toolbar attachment manager could  
25 read it from the EDF file, instead of from a URL. The information contained in

1 the EDF is processed and (together with the information contained in the PKG),  
2 the appropriate files are automatically installed on a user's computer. This is done  
3 unobtrusively without manipulating the computer's persisted settings, as might be  
4 found in the user's system registry.

5 An EDF, implemented in XML, contains various tags that are associated  
6 with various extensions. The various tags can correspond to:

- 7 • User interface elements
- 8 • Behaviors/Components/Objects
- 9 • Store Elements
- 10 • User-defined objects
- 11 • Or anything else that represents a point of extensibility in the  
12 application or platform

13 EDFs advantageously have an "open schema" which means that third party  
14 developers can extend the extension mechanism and include their own extensions  
15 by creating their own tags. Additionally, extensions can themselves be extended  
16 by other developers. EDFs can also have one or more predefined tags. Exemplary  
17 predefined XML tags for user interface elements can include tags for feature types  
18 such as: tool bars, accelerators, menu items, and themes. These feature types are  
19 utilized in the single navigable window application incorporated by reference  
20 above and defined in the table immediately below:  
21  
22  
23  
24  
25



Feature Type	Definition
Tool Bars	Horizontal command containers above the document area.
Accelerators	Keyboard shortcuts for commands
Menu Items	Pop-up or drop-down menu choices that third parties can add to well-known, named menu attachments in the platform
Themes	A data-driven way to provide overrides for well-known resources of the platform, such as default buttons or default style sheet

Table 1

Exemplary predefined XML tags for behaviors/components/objects include tags for Services. These feature types are utilized in the single navigable window application incorporated by reference above and defined in the table immediately below:

Feature Type	Definition
Services	Services are objects that extend existing objects (such as the application, window, or document) in the application or platform's Object Model. For example, editing functions use Object Model Attachments attached to the window or document that maintain document context and editing state per-window. These can also include Object Model Attachments attached to the application (such as a spellchecker dictionary object)

Table 2

Exemplary predefined XML tags for store elements include tags for content classes and offline data sources. These feature types are utilized in the single navigable window application incorporated by reference above and defined in the table immediately below:

Feature Type	Definition
Content Classes	Allow extension writers to define new types of XML documents with new schemas.
Offline Data Sources	Allow for extension writers to define store replication instructions in an EDF.

Table 3

### EDF Schema

In the described embodiment, the EDFs have a particular XML schema that is utilized. The schema comprises collections of XML tags that are arranged in a hierarchical organization to facilitate information dissemination to software components that need certain extensions. In the described embodiment, the outer (encompassing tag) for EDFs is an “extension” tag.

Fig. 4 shows an exemplary extension tag. “extension” tags can include one or more of the following attributes, all of which are optional:

Attribute	Definition
urn	ID for the extension. It allows extension writers to specify relative locations for content in EDFs without using relative paths or fixed URLs. It also allows hosting administrators to move around extensions on servers without breaking any links.
name	Name that can be used in a status bar or message display
version	Vendor-determined version number for the extension.
lastUpdate	Date/time that the EDF was last modified.
description	Brief description of the extension.

Table 4

Within the “extension” outer tag are one or more child tags, also referred to as “top level tags”. These top level tags are each associated with a feature type

that can be added by a particular extension. Exemplary feature types are discussed in connection with Tables 1-3 above. Underneath each top level tag there can be one or more child tags that are individually associated with a particular feature of the feature type that is to be added by a particular extension.

Fig. 5 shows an exemplary XML schema organization in accordance with this embodiment. For each top level tag in an EDF, there is an associated attachment manager which is a software component that receives data associated with the tag so that the data can be used to incorporate the extension into the platform or application program. Different attachment managers may interpret the data from the tag in different ways to provide different types of extensibility, so different top level tags will contain different types of data in different structures. This will become more evident in the "Architecture" section below. Note that the "edf:" XML namespace qualifier enables support of an open schema where extensions can provide their own tags and corresponding attachment managers. Tags within the edf namespace are used by built-in attachment managers in the application or software platform. Tags in other namespaces are used by third-parties to provide additional points of extensibility.

### **Package Manifest (PKG file)**

The package manifests (PKGs) assist in organizing the downloading of software in the form of multiple files over a network such as the Internet. The PKGs are advantageously employed, in the example given in this document, with EDFs. As pointed out above, however, the techniques discussed in connection with the PKGs can be deployed independently of EDFs and in any suitable scenario where it is desirable to deliver software over a network such as the

1 Internet. While the EDFs describe the logical attachment of extensions into an  
2 application program or platform, the package manifest's role is to assist in one or  
3 more of: organized delivery, validation and/or, updating of files associated with the  
4 various extensions that can be provided.

5 In designing a delivery mechanism for Web-assisted delivery of software  
6 content or files, several considerations are of interest.

7 Whenever possible, it is desirable to reduce the size of required downloads  
8 during update and installation operations. To address this consideration, software  
9 content is broken into multiple packages. Each package contains a group of one or  
10 more files that implement a common or well-defined functionality. By breaking  
11 the content into individual packages, the size of the required download during  
12 installation and update can be minimized. Each package is then described by a  
13 package manifest (PKG) that includes file information such as file locations and  
14 hash values that can be used for validation or security and versioning.

15 It is also desirable to give the end user a Web-like experience. To do this,  
16 extensions are loaded in a manner that feels to a user more like they are loading a  
17 web page, rather than traditional software packages where the user has to wait  
18 until the entire package is loaded before they can interact with it. In the described  
19 embodiment, users are given a web-like experience by streaming the extension  
20 files down to the client so that a user can begin to interact with an application  
21 program much sooner than if they had to wait for the entire software application  
22 program to load. For example, if there are user interface (UI) image files  
23 streaming down, the user can see the UI as the files stream in. Consider, for  
24 example, the single application program having the multiple different  
25 functionalities that is described in the patent application incorporated by reference

above. A user might browse to an email functionality and download the files that are necessary to interact with the email functionality. Files that are associated with another different functionality would then be downloaded after the files associated with the email functionality. In this way, the user can begin to operate within a particular functionality without having to wait for all of the files associated with all of the other functionalities.

Another consideration of interest pertains to the efficiency with which the extension files or “bits” are delivered to the client. To address this consideration, the described embodiment utilizes a couple of different download approaches: a throttled download and a background download. Throttled downloading conducts download operations while taking into account the available bandwidth and type of media over which the files are being transferred. Any suitable throttled download process can be employed and will be understood by those of skill in the art. Background download is conducted while the user is working within the application program and is implemented by allocating a background thread so that the user can continue their work. One optimization that is utilized is that packages are prioritized and delivered in accordance with what a user might be working on.

Another consideration is associated with optimizing the user's computing experience. Here, the user's experience is optimized by making available the most common scenarios for the user. This is effectively directed to giving the user the functionality that they want first, and then, through the background download process, providing the code that implements functionalities that the user might use in the future. To determine which functionalities a user desires to have first, an automated scenario-based packaging process is provided that runs against file usage logs from scripted scenarios.

1 All of these considerations and the inventive solutions directed to  
2 addressing the considerations are discussed in detail in the sections that follow  
3 below.

#### 4 5 Package Manifest Definition

6 In the described embodiment, a package manifest (PKG file) comprises a  
7 list of files that are utilized in a package. The list is advantageously compressed  
8 somewhat and digitally signed. Each package manifest can contain a list of one or  
9 more files each of which can include an associated hash, as well as download  
10 directives that control caching of the files. Once an extension is authored, a  
11 software tool can be used to generate the package manifest.

12 In addition, the package manifest can specify several other pieces of  
13 information:

#### 14 15 • FILE GROUP

16 All files in an extension can be labeled according to a number of predefined  
17 file groups. The file group of a particular file determines when the particular file  
18 gets downloaded, where it is stored on the client, and how it gets packaged. In the  
19 described embodiment, four predefined file groups are provided and are listed and  
20 described in the table immediately below:

21  
22

23 Group name	When downloaded	Where stored on the client	Packaging	Content
24 Required	Downloaded before any other files in the extension.	NetDocs package cache	All required files in an extension are	DLLs included so that a user will not have to wait for a prolonged

25

			packaged together as a CAB* file.	period of time before clicking on a UI element
Offline	Offline files start getting downloaded as soon as Required are down. Providing the user stays on line long enough, these files will all get downloaded and will later be available for offline use.	NetDocs package cache	File are sent down individually.	Bulk of the UI files.
On demand	Only downloaded when they are requested for the first time.	NetDocs package cache	Files are sent down individually.	To avoid using up disk space on the client, advanced features can be put in this category.
Online only	Downloaded on demand. Content is only available when the user is online.	WinInet Cache	Files are sent down individually	Content that is not to be provided offline. Examples include help pages and other content that can consume a large amount of disk space.

\* CAB stands for the CABinet technology that Internet Explorer uses to package bits for download. CAB files average from 2 to 3:1 compression, and are optimized for quick expansion. CAB files have the added security benefit that they are easily signed.

## • FILE DOWNLOAD PRIORITY

Files in each group are listed according to the order in which they should be downloaded. This download order is implicit in the ordering of the files in the package manifest, an example of which is shown in Fig. 6.

## • HASH VALUE FOR SECURITY/VERSIONING

Individual files in the package manifest can have an associated hash value. Each hash value is generated by running the file through an encryption algorithm. An exemplary encryption algorithm is Microsoft's CryptoAPI. In the illustrated example, each file can be listed with a base 64-encoded hash value, so that the file can be validated once the content arrives at the client. Specifically the package manifest is sent to the client in a secure manner (i.e. it is digitally signed). The

1 package manifest contains the hash values for individual files. When the  
2 individual files are received by the client, each of the files can be run through the  
3 same Crypto API that was used to provide the hash values in the package  
4 manifest. If the hash values for a particular file compare favorably, then the file  
5 has not been altered and is secure.

6 When a file is updated, hash values can serve a useful purpose in  
7 identifying files that have not been changed between different versions of an  
8 extension. Consider Fig. 7, for example. There, an old directory 700 in a client  
9 package cache contains package A which include two files—file 1 with hash = x,  
10 and file 2 with hash = y. Assume that this package is associated with an older  
11 version of an extension. When an updated version is produced, its package  
12 manifest is delivered to the client. The updated extension version is represented  
13 on a source directory of a code or web server 704. The package manifest includes  
14 the hash values for all of the files in the new extension version. A new client  
15 destination directory 702 is defined for all of the files of the new extension. If any  
16 of the hash values for files in the new extension are the same as the hash values of  
17 the files in the old directory 700, then those files can be copied directly from the  
18 old directory 700 to the new destination directory 702. In this example, file 1's  
19 hash value is the same as the hash value for file 1 on the source directory 704, so it  
20 can be copied into the new destination directory 702. File 2's hash value, however  
21 is different from the hash value for file 2 on the source directory, so it is not copied  
22 from the old directory 700. Rather, file 2 is downloaded from the code server. A  
23 new file 3 has been added and is also downloaded from the code server. Hence, in  
24 this example, a new version of an extension resulted in a download of less than all  
25 of the files in the extension version. This is because hash values for each of the





1           • DLL LOAD DEPENDENCIES

2           The reason for the dependencies section is to allow for legacy code that  
3 relies on being loaded by virtue of being in the search path of some other dll. In  
4 this case we have to make sure that the dependency dll is in the package cache  
5 directory before the dependant dll is loaded. Fig. 6 shows an exemplary package  
6 manifest 600 that is defined in a hierarchical tag-based language. Advantageously,  
7 the tag-based language comprises XML which is desirably extensible and flexible.  
8 In this example, a number of tags are provided in a hierarchical arrangement. The  
9 “package” tag contains information about the size of the package. The “files” tag  
10 is a child of the “package” tag and contains information about the file groups that  
11 are contained in that particular package. The “file” tag is a child of the “group”  
12 tag and contains information about specific files that comprise the extension, i.e.  
13 file name and hash value. A “dependency” tag is provided as a child of the “file”  
14 tag and lists any dependencies as discussed above. A “COMClass” tag is also  
15 provided as a child of the “file” tag and contains IDs as mentioned above. The  
16 ordering of the file groups in this schema implicitly defines the download order of  
17 the files.

18  
19           Package Delivery

20           To optimize package delivery, two different delivery schemes are utilized.  
21 First, a throttled download approach is utilized using known throttling download  
22 techniques. Here, considerations such as available bandwidth and media over  
23 which the extensions are being provided are considered.

24           Second, a background download approach is utilized. Background  
25 downloads enable a user to continue to work within an application program while

content is downloaded. Foreground downloads are used when the user has explicitly requested a file/extension by clicking, for example, on an extension link, or requested an action, for example, by clicking on the “Compose” mail button, that requires download of files which are not available locally.

Along with background downloads, a queue management feature is provided. Specifically, when an extension must be installed or updated, a package manager, which is essentially a software component that manages packages, is provided with the following information:

- URL for the package manifest information on a code server
- URN for package destination directory in the package cache at the client
- (Optional) URN for the old package directory (if one exists) in the package cache

From this information, the package manager creates a package object and adds the package object to a download queue. The download queue is designed for easy rearrangement of a package download order. Consider, for example, Fig. 8 which shows a portion of a download queue 800 that contains two package objects—package object 802 (corresponding to package A) and package object 804 (corresponding to package B). The package objects maintain a list of which files of a corresponding package have been downloaded or installed. In the present example, files 1 and 2 from package A have been installed while file 3 has not been installed; and files 1, 2, and 3 have not been installed from package B. The download queue can be rearranged based on what the user is doing. That is, based on the actions that a user takes, the priority of files that are to be downloaded can be changed. In this example, the package manager is designed to process the first uninstalled file in the package at the head of the download queue.

1 If, however, the user starts to use a file in an extension that is different from the  
2 extension whose files are at the head of the download queue, the corresponding  
3 package for the file that the user has started to use can be moved to the head of the  
4 download queue. Because a file's package is specified by its URN, the file's  
5 package can be quickly identified and located in the download queue. For  
6 example, and considering Fig. 8, if one of the files in package B is requested  
7 before the package manager has started to install package A's third file, then  
8 package B will be moved to the head of the download queue.

9 Fig. 9 is a flow diagram that describes steps in a download queue  
10 management method in accordance with the described example. The method can  
11 be implemented in any suitable hardware, software, firmware or combination  
12 thereof. In the present example, the method is implemented in software.

13 Step 900 receives one or more requests for an extension. The requests can  
14 be generated in any suitable way. Step 902 creates a package object that  
15 corresponds to each extension package that is to be downloaded. Step 904  
16 arranges the package objects in a download queue. Step 906 then downloads files  
17 corresponding to the package objects in the download queue. This step can be  
18 implemented by, for example, starting at the head of the download queue and  
19 downloading files until all of the files for a package object have been downloaded,  
20 and then moving to the next package object. Step 908 ascertains whether a user  
21 action requires a file that is not described in the current package object. If the  
22 user's action does not require a file not described by the current package object,  
23 then the method branches back to step 906 and continues to download files  
24 associated with the current package object. If, on the other hand, the user's action  
25 requires a files that is not described in the current package object, then step 910

007250-2786666

1 moves the package object associated with the required file to the head of the  
2 download queue and begins to download files associated with this newly-  
3 repositioned package object. This step can be implemented by ascertaining which  
4 package object is associated with the required file by ascertaining the URN  
5 associated with the file. This URN specifies the file's package so that its package  
6 object can be quickly located and moved to the front of the download queue.

### 7 8 Package Creation

9 One of the innovative features of the described embodiment is its  
10 extensibility. That is, a software platform is provided in the form of an application  
11 program that can be extended by various third-party user-defined extensions.  
12 These extensions are delivered via the Web and are integrated directly into the  
13 software platform. In order to provide an organized delivery process, packages  
14 should be created in a uniform manner so that they can be predictably handled and  
15 integrated into the software platform.

16 In accordance with the described embodiment, each package should  
17 correspond to an end-user feature. For example, in the patent application  
18 incorporated by reference above, separate packages are provided for each of the  
19 email, contacts, document authoring, and planner functionalities. If packages that  
20 do not depend on one another share a dependency, then this shared dependency  
21 should become its own package. For example, there is no reason why the email  
22 and document authoring functionalities should depend on one another, yet both of  
23 them require the ability to publish content. By separating the publishing  
24 functionality into its own package, a certain amount of download order flexibility  
25 is preserved. Depending on what the user starts to work on, the files

1 corresponding to the email functionality or the document authoring can be  
2 downloaded first.

3 Fig. 10 is a flow diagram that describes steps in a package creation method  
4 in accordance with the described example. The method can be implemented in  
5 any suitable hardware, software, firmware or combination thereof. Portions of the  
6 method might, however, be implemented manually.

7 Step 1000 identifies end user features that are to be provided as extensions.  
8 Step 1002 identifies any shared dependencies among the end user features. Step  
9 1004 creates individual packages for the end user features. Step 1006 creates  
10 individual packages for any shared dependencies among the end user features.

#### 11 12 Automated Package Manifest Creation Tool

13 Advantageously, and in accordance with one implementation, an automated  
14 package manifest tool is provided and takes various input parameters and  
15 automatically creates a package manifest. The tool can be available to third  
16 parties to assist them in creating a package manifest.

17 Fig. 11 shows an exemplary package manifest creation tool 1100 that is  
18 desirably implemented in software. In this specific example, the tool can take the  
19 following input parameters (some of which are optional):

- 20
- 21 • Extension directory
  - 22 • File group information and DLL load dependencies (Optional)
  - 23 • File usage statistics from scenario runs (Optional)

24 The extension directory input parameter specifies the directory containing  
25 all of the files that will be described by the package manifest. If this is the only

parameter, then tool 1100 will generate a manifest in which the EDF and DLLs in the directory are listed in the "Required" set, and all other content is "Offline".

The file group information and load dependencies parameter is optional. If an extension author has an idea of the categories in which his or her files should be placed, the categories should be specified here. For example, the author of the template manifest shown below knows that he wants his error handling GIF to be included in the package's required set. His choices here will always be respected in the final manifest. Additionally, if the extension author knows of any DLL load dependencies, these should be specified here as well.

```
<?xml version="1.0"?>
<Package>
  <Files>
    <Group Name="required">
      <File Name="bar.dll"/>
      <File Name="foo.dll"/>
      <Dependencies>
        <File Name="bar.dll"/>
      </Dependencies>
    </File>
    <File Name="myextension.edf"/>
    <File Name="errorhandling.gif">
  </Group> ...
</Files>
</Package>
```

The file usage statistics from scenario runs parameter is an optional parameter. This parameter enables the file download priority to be determined based on scenario runs. A scenario is a script of tasks that the average user typically follows when using a product during a particular portion of product use. For example, one scenario might pertain to the tasks involved in sending an email

message (i.e. click “new mail” button, type in “TO” well, type is “Subject” well, etc.). In the described embodiment, file usage statistics from scenario runs are collected from running IIS logs on various scenarios. The different scenarios are directed to ensuring, with some degree of probabilistic support, that the file download order reflects, in some way, the files that will likely be used by the user first.

It will be appreciated that the file usage statistics can be provided dynamically by building a knowledge base that describes the actual tasks that people typically accomplish. The information that is maintained in the knowledge base can then be used to generate and adapt download scenarios that actually conform to patterns that are established across a user base.

If extension writers want to pass the package manifest creation tool 1100 this information, they need to specify the log directory, as well as the start and end dates of the section of the log that the tool should analyze. For third parties, the download priority order within a group will be the order in which the group's files were requested in the logs across all scenarios.

In one implementation, the approach is somewhat more sophisticated. Additional information (in addition to the scripted steps) is stored in the IIS logs and includes scenario priority and checkpoints. The scenario priority is a priority that is assigned for each scenario. So, for example, if one scenario is ten times more important than another scenario, this information can be maintained. The priority (e.g. a rating from between 1 to 100, with 100 being the highest priority), should be equal to a best guess as to the percentage of the time that users will step through the scenario, assuming they use the extension at all. Checkpoints provide a way to separate one scenario from another. For example, checkpoints designated



1 as "Offline" and "Shutdown" can be automatically added at the beginning and end  
2 of the scenarios, respectively, so there can be a differentiation between scenario  
3 runs in the log. Additionally, script authors can optionally use checkpoints mid-  
4 scenario to indicate a change in group priority, e.g. part of a scenario script could  
5 be labeled as an "On demand" feature and another part could be labeled as  
6 "Offline".

7 Fig. 12 is a flow diagram that describes steps in a package manifest creation  
8 method in accordance with the described example. The method can be  
9 implemented in any suitable hardware, software, firmware or combination thereof.  
10 In the described example, various steps of the method are implemented by a  
11 manifest creation tool that is implemented in software.

12 Step 1200 provides a package manifest creation tool. This tool can be a  
13 software tool that resides on an extension author's machine. Step 1202 receives  
14 information pertaining to an extension directory as a first input parameter. Step  
15 determines whether there is any file group information or load dependency  
16 information provided by the extension author. If there is, step 1206 receives the  
17 information as an input parameter. Step 1208 determines whether there is any file  
18 usage statistic information. Such information can be provided, in one  
19 embodiment, through the use of scenario runs as described above. If such  
20 information is provided, then step 1210 receives the information as an input  
21 parameter. Step 1212 then uses all of the information provided as input  
22 parameters to generate a manifest automatically.

23  
24  
25

## Exemplary File Ordering Heuristics Based on File Usage Statistics

Fig. 13 is a flow diagram that describes but one exemplary file ordering or sorting heuristic in accordance with the described embodiment. It is to be understood that this specific example constitutes but one way of ordering files for download. Accordingly, other heuristics can be used without departing from the spirit and scope of the claimed subject matter.

Step 1300 sorts files by file group. Recall that in the illustrated example above, files can be grouped in one of four possible groups: Required, Offline, On Demand and Online Only. A file's group is determined first by the manifest, and, if it does not provide any group information, then by the highest priority group that it uses, according to checkpoint information in the log. Files in the "Required" set should not be considered because their order is already known. If no group information is included about a file, then an assumption is made that the EDF and all DLLs are "Required" files and all other files in the directory are "Offline".

Consider, for example, the following initial file usage information for three different scenarios:

Scenario 1 file usage: 1) FileA.gif, 2)FileB.xml, 3)FileE.dll

Scenario 2 file usage: 1) FileC.xml, 2) FileA.gif

Scenario 3 file usage: 1)File D.js, 2)FileA.gif

Scenario 1 = priority 80

Scenario 2 = priority 80

Scenario 3 = priority 40

In this example, there are three scenarios that have files associated with them. Each of the scenarios has a priority with which it is associated. The files

are first sorted by group (step 1300). Recall that in this ordering heuristic, DLLs are "Required" and all other files are considered "Offline". This provides the following sorted files:

Required files  
FileE

Offline files  
FileA, FileB, FileC, File D

Step 1302 sorts files based on scenario priorities (from highest to lowest). Higher priority files are ordered so that they are downloaded first. This step provides the following sorted files:

Required files  
FileE

Offline files  
Priority 80 group: files used by Scenarios 1 & 2 = File A, File B, and File C

Priority 40 group: files used by Scenario 3 (that are not already listed) = File D.

Step 1304 then sorts the files by file usage order within a scenario run. For each priority grouping with more than one file, the files are sorted according to the average order in which they were downloaded within scenarios of their labeled priority. Scenarios with a smaller average usage order will be downloaded earlier. Ties are broken based on the order in which the scenarios appear in the input file. As an example, consider the following:

File A: average order = (Scenario 1 order + Scenario 2 order)/2 = (1+2)/2 = 1.5.

File B: average order = (Scenario 1 order)/1 = (2)/1 = 2.

File C: average order = (Scenario 2 order)/1 = (1)/1 = 1.

Here, file A got used first by scenario 1 and second by scenario 2 for an average of 1.5, and so on. File C has the smallest order number so, of the Offline files, it is sent first. The final file order is shown below:

Required files

FileE

Offline files

FileC, FileA, FileB, File D

**Code, Components and “bits”**

The following files and resources can be, but need not be included with an extension. This list is not exclusive, as other resources can certainly be incorporated into an extension.

- Customized UI and keyboard shortcuts
- Components and Behaviors
- XML browsing and editing components (including XSL and business logic objects)
- Static pages or other resources
- Third-party defined custom content

Users install extensions by navigating to a network site for the extension. In an Internet implementation, the user navigates to an appropriate URL for the extension. Hosting administrators can also “push” extensions so that users can

1 automatically receive them by adding an entry into the appropriate users'  
2 "Preference" settings.

### 3 4 **Platform Set Up and Extension Installation**

5 Fig. 14 is a flow diagram that describes exemplary steps in a set up and  
6 extension installation process in accordance with the described embodiment. This  
7 example describes an Internet-based example. In the illustrated example, various  
8 extensions are maintained at or accessible through various Internet sites. The  
9 extensions are deliverable via the Internet to a client. It will be appreciated that  
10 the illustrated division of computers may not necessarily exist at all. For example,  
11 all of the functionality embodied by the computers may reside on one machine, or  
12 the extension may be local, or the platform and the extensions may be on the same  
13 machine, etc.

14 The flow diagram, in this example, is illustrated in connection with three  
15 separate "zones", one of which represents a client, one of which represents a  
16 "platform" Internet server, and one of which represents a third party Internet  
17 server. The acts that are described in connection with the different zones are  
18 performed, in this example, by the entities assigned to the zone. In some  
19 configurations, one or more of these zones may overlap. For instance, the  
20 platform server may be the same device as the extension server.

21 Step 1400 navigates a user to a particular Internet site that is associated  
22 with the software platform that is to be utilized as the foundation for extension  
23 installation described below. In step 1402, the user clicks an "install" button that  
24 sends a message to the software platform server that indicates that a user wishes to  
25 install the software platform. This step can be an optional step. Steps 1404 and

The steps described immediately above constitute steps associated with an initial set up in which the software code for the single navigable window application is delivered to and installed on a client machine. The steps described immediately below are associated with extension installation.

Step 1410 uses a link that is associated with an extension to access the extension. This step can be implemented by a user navigating their browser to a particular Internet site through which one or more extensions can be accessed. Alternately, a reference to the link can be placed in the user's preferences or the preferences of a computing group with which the user is associated (e.g. the system administrator can place a reference in a group's preferences). The link can advantageously be associated with a third party Internet server or Internet site. Step 1412 downloads extension files according to the PKG associated with an EDF. The files are delivered to the client and step 1414 places the extension files in a local store as specified by the PKG specification. At this point, an extension is installed and the user can utilize the functionality provided by the extension. Step 1416 determines whether extension updates are available. This can be done by periodically polling an extension catalog (discussed in an "Extension Catalog"

section below) to ascertain whether there are any updates to various extensions. Alternately, notifications might be automatically sent to the client so that the client is aware of updates or any other method might be used to determine if updates are available. If there are updates available, step 1418 branches to step 1412 which downloads the extension files associated with the update and installs them on the client.

### **Developing Extensions**

Developing extensions for the software platform is a fairly straight-forward process. A developer develops the extension content using a tool such as Notepad or other tools such as Visual Studio. The extension is then described in an EDF and PKG and the PKG is digitally-signed and then optionally compressed. The extension can then be hosted on a particular network server.

Fig. 15 is a flow diagram that describes steps in an extension development method in accordance with the described embodiment. One or more of these steps can be performed by a software developer or organization that creates a particular extension. Some of the steps are implemented in software. Step 1500 develops an extension. Any suitable tools can be used to develop the extension. Step 1502 then creates an extension definition file (EDF) for the extension. The EDF, in this example, is defined using XML as discussed above. Other formats can, of course, be used to describe the EDF. Step 1504 creates a package manifest (PKG) for the extension. The PKG, in this example, is defined using XML as discussed above. Step 1506 then hosts the EDF and PKG on a network server such as an Internet server. Additionally, the associated extension files that are described in the PKG can also be hosted by the network or Internet server (step 1508). Having

accomplished the above, users can now navigate to an EDF directly (using, for example, the associated URL or some other network address), which then installs the extension by caching any required files locally and placing a reference to the extension in the user's preferences.

Specifically, step 1510 delivers the EDF and PKG files to a client. This step can be implemented by a user navigating to a particular Internet site where the appropriate files are hosted and then downloading the files. Step 1512 delivers the extension files that are associated with the EDF and PKG files to the client, whereupon they can be installed and used.

### **Extension Catalog**

One optimization, discussed briefly in connection with Fig. 14, is an extension or EDF catalog which provides an additional level of indirection to the EDF. An EDF catalog allows organizations to group extensions and provides a single place to determine when an extension changes. The desired extension can be automatically selected from the catalog by the software platform based upon the user's settings. The catalog can be queried to determine which extension is most appropriate for the user.

In the described embodiment, a catalog is an XML file that contains mappings from extension URNs to one or more package URNs based upon language, version or other attributes. Catalogs can, however, be defined using any suitable format. Catalogs can provide:

- The ability for a hosting organization to update version information for one or more hosted extensions in a single place



- Optional automatic indirection to correct version based upon user's settings. For instance, a catalog may list several versions of an extension for different languages. The catalog file can be processed to find the version of the extension that matches the user's language settings.
- Optional automatic upgrade to new versions of extensions as they become available

Like EDFs, catalogs can be compressed and digitally signed to prevent tampering. By subscribing to a catalog in order to detect changes for one or more hosted extensions, the number of server pings required from the client (or notifications to the client) in order to discover extension updates can be reduced.

Fig. 16 shows an exemplary XML catalog structure. Entries in the catalog can be structured as follows:

Attribute	Type	Required	Description
extensionURN	uri	Y	Identifier for an extension. There may be more than one entry for a given extension urn in a single catalog representing different versions, languages, etc.
name	String	N	Friendly name for an extension.
packageURN	uri	Y	Required urn for the package. Package urn corresponds to a discrete set of bits. It is different from the extension urn: For each user, the extension urn (name) corresponds to a specific set of files based upon language preferences and version. This means that, for shared machines, different users may have different extensionURN to packageURN maps based

			upon their preferences.
packageURL	uri	Y	url of the digitally signed compressed file containing the PKG file is required.
language	String	N	Language is an optional language specifier.
version	String	N	Version is an optional version specifier
defaultLanguage	String	N	DefaultLanguage is an optional attribute specifying the default language package. For a given version of an extension, there should be only one entry with DefaultLanguage attribute.
defaultVersion	String	N	DefaultVersion is an optional attribute specifying default version for an extension. For a given extension urn and language attribute there should be only one entry with DefaultVersion attribute.

In this particular example:

- The default language of the netdocs-planner is the English version.
- The default English version is 1.1. The default French version is 1.0. If there is no version available in the user's specified language on the platform, they will get the English version 1.1 by default.
- The English version of netdocs-planner has been upgraded from V1 to V1.1.
- There is also a French version. The extension URN is the same as the English version. There is no 1.1 release for French yet, so 1.0 is the current version for French speaking users.
- A query against the catalog returns only the rows where language matches the user's language preferences. The query would also return all rows where language is the user's language or default = 'yes' and throw out duplicates for the same name.

## Architecture

In the described embodiment, one of the aspects that provide desirable utility is the extensibility of the software platform. That is, third and fourth party developers are free to develop their own extensions which can then be used within the framework of the software platform. The extensions are integrated directly into the software so that the platform's functionality is modified by the extensions. Recall that in order to provide an extension, the developer simply authors the extension, describes their extension in an EDF and PKG, and then hosts the EDF, PKG, and associated files on a network server.

The EDF, as pointed out above, can be defined in an XML schema that includes a root node (i.e. the "extension" tag) and one or more child nodes. In this particular example, the child nodes generally correspond to the individual extension feature types that are desired for incorporation into the software platform. For example, recall that Tables 1-3 above describe various exemplary predefined feature types that can be added through an extension using the predefined XML schema.

Consider now a developer who wants to add two menus and a toolbar to the software platform. The menus and toolbar might be associated with a retail store that maintains a Web site for its customers. The retail store might want a customer who visits its Web site to be presented with a UI that is unique to the retail store and provides services that are specifically tailored to the store. To do this, the developer develops two different menus, one of which might be associated with displaying the most recent specials, and other of which might be associated with providing a search mechanism through which the user can search for specific products. The toolbar might contain specific buttons that are unique to the retail

store. A simplified EDF called “*retail.edf*” for the retail store’s extension is shown directly below:

```
<edf:extension name= “retail extension” urn= “extension.retail.com”>
  <edf:menus>
    <edf:menu url= “url1.htm”/>
    <edf:menu url= “url2.htm”/>
  </edf:menus>
  <edf:toolbars>
    <edf:toolbar url= “url3.htm”/>
  </edf:toolbars>
</edf:/extension>
```

Here, the outer “extension” tag designates this XML file as an extension. The inner “menus” and “toolbars” tags are top level tags that designate that the information between these tags pertains respectively to menus and toolbars that correspond to the extensions that the developer has added. The boldface “menu” and “toolbar” tags describe data pertaining to the actual extension and contain a URL that is associated with each extension as described above. The EDF above logically describes the extensions that are being provided as including two menus and one tool bar.

Consider also that the above EDF can constitute but one of many EDFs that are loaded into the system. Each EDF can contain one or more top level tags, each of which is associated with one or more specific extensions that are to be added to the software platform.

Fig. 17 is a block diagram of an exemplary software architecture that is configured to process multiple different EDFs so that the software components that are responsible for incorporating each particular extension into the software

1 platform receive the appropriate information that is specific to their extension.  
2 This example is specific to the XML implementation that is discussed throughout  
3 this document. It is to be understood that other architectures, similar in  
4 functionality to the one discussed below, can be used in other implementations  
5 without departing from the spirit and scope of the claimed subject matter.

6 Utility objects, herein referred to as attachment points, are used to process  
7 the information from the multiple EDFs. An attachment point is simply a  
8 collection of objects that fire events to registered listeners as objects are added to  
9 or removed from the collection. Many types of attachment points can be created,  
10 but all take data from a source (often another attachment point), process it (either  
11 dynamically or statically), and expose the results of their processing. Some of the  
12 simplest attachment points include:

- 13 • An XML attachment point, which loads an XML file and exposes  
the top-level nodes of the XML as objects in its collection.
- 14 • A filter attachment point, that connects to another attachment point  
and exposes only those objects from it that meet some criteria.
- 15 • A merge attachment point, that connects to one or more other  
16 attachment points and exposes all of their objects as one, merged  
collection of objects.

17 In the illustrated example, the architecture includes a collection of one or  
18 more attachment points, including a funnel structure known as an EDFHub 1700,  
19 an attachment point manager 1702, and multiple attachment managers 1704. The  
20 EDFHub 1700 receives all of the EDFs and merges them together and exposes  
21 them as a single list. Other individual attachment points provide mechanisms that  
22 manipulate (including filter, merge and expand) the single list that is exposed by  
23 the EDFHub 1700. Whenever a new extension or EDF is added to or removed  
24 from the EDFHub, the various attachment points will see to it that the appropriate  
25 attachment manager(s) is notified. This is done by firing events to the appropriate

attachment managers. The attachment point manager 1702 creates, destroys and manages the various attachment points in the system and allows easy reuse of identical attachment points.

For each top level tag (i.e. "menus" and "toolbars" tags), there is a corresponding attachment manager 1704 that uses data provided by the attachment points to incorporate a particular type of feature within the software platform. Each attachment manager requests a set of attachment points from the attachment point manager 1702. These manipulate the data exposed by the EDFHub 1700. In the illustrated example, the attachment points can be requested as a predicate chain that the attachment point manager uses to create and build a set of attachment points that operate on the data exposed by the EDFHub 1700.

Fig. 18 is a flow diagram that describes steps in a method in accordance with the described embodiment. The method is implemented in software and, in this example, by the software components of Fig. 17.

Step 1800 receives multiple EDFs. These files can be received in any suitable way. For example, a user can specify in their preferences particular extensions that they desire to be loaded when they are online. Alternately, a user might navigate using a link to a particular Internet site that recognizes that the user is executing a software platform that is configured to dynamically add the extensions. The EDFs, in this example, are funneled into the EDFHub 1700 which uses attachment points to combine the EDFs (step 1802). In this example, the EDFs are defined as XML files and the nodes are combined into one single XML list. Step 1804 exposes the combined EDFs. In this particular example, the EDFs are combined into a single XML list that is exposed to other various attachment points which then further manipulate that data (step 1806). One goal of the

attachment points is to avoid having the attachment managers 1704 re-query the whole system every time an extension gets added or removed from the system. Thus, if an extension gets added or removed, the attachment points ensure that only the appropriate attachment manager 1704 is notified of the specific additions or deletions of the extension. For example, if an EDF indicates the addition of a menu, then only the attachment manager associated with menus is notified. Accordingly, step 1808 notifies the appropriate attachment manager of any new data that matches the attachment managers requirements.

### **Attachment Points and the Attachment Point Manager**

Attachment points are objects that expose collections of ordered objects and fire notifications when new objects are inserted or deleted. In the exemplary system, the objects are XML nodes, but they could be any type of object. Although there are many types of attachment points, they all follow a similar process:

- 1) Initially attach to one or more data sources. These could be files or, commonly, other attachment points.
- 2) Process the data based on some logic. Usually the logic is quite simple and could involve something like filtering the objects based on some criteria.
- 3) Expose the results of the processing step 2 in a new collection of objects.
- 4) Fire events to indicate how the exposed collection of objects changed (OnInserted(index, count) or OnRemoved(index, count)).
- 5) Optionally, continue to listen for changes in the data sources and repeat step 2-4 when changes occur.

Alone, each attachment point is quite simple, but when different types of attachment points are combined to form "chains", where one attachment point processes data from a second attachment point, the processing can be quite





1 node of the XML file as the only object in its collection. The inner Explode  
2 attachment point uses the URL attachment point as its data source and exposes all  
3 of the children of the objects in that source collection. In this case, the children of  
4 the root node are the top level XML tags “menu” and “toolbars”. The Filter  
5 attachment point uses the Explode attachment point as its data source and filters  
6 the exposed objects looking only for the nodes that are “menus”. The outer  
7 Explode attachment point uses the Filter attachment point as its data source and  
8 exposes all of the children of the filtered menu nodes to provide a list containing  
9 the two menus that are added by the extension. Since this particular XML file  
10 contained menus that were identified by the attachment points associated with  
11 menu attachment manager, that attachment manager is then notified that two  
12 menus have been added by an extension.

13 This process is diagrammatically illustrated in Fig. 19 which shows  
14 attachment points 1900, 1902, 1904, and 1906. Each attachment point exposes a  
15 list of XML nodes. URL attachment point 1900 takes an input (a URL to an XML  
16 file—e.g. *retail.edf*) and exposes a list of XML nodes. This list contains only the  
17 root node “<edf:extension>”. Explode attachment point 1902 takes as an input  
18 attachment point 1900 and exposes a list of XML nodes which are children of  
19 source XML nodes. In this example, the list of XML nodes exposed by  
20 attachment point 1902 are the “<menus>” nodes and the “<toolbars>” nodes. The  
21 filter attachment point 1904 takes attachment point 1902 as an input and filters on  
22 “menus.” It then exposes an XML list having only “<menus>” nodes in it. The  
23 explode attachment point 1906 takes attachment point 1904 as an input and  
24 exposes a list with the XML nodes that are contained in the “<menus>” nodes—  
25 here both of the “<menu>” nodes.

1       Consider additionally that the toolbar attachment manager would request a  
2 predicate chain of attachment points which would also use URL attachment point,  
3 an Explode attachment point and a filter attachment point 1904 that filters on  
4 "toolbars". Thus, the corresponding explode attachment point 1906 would expose  
5 an XML list containing only the "<toolbar>" node. But, the attachment point  
6 manager would detect the commonality of the URL attachment point and the inner  
7 Explode attachment point, so it would reuse the same attachment points it created  
8 for the menu attachment manager. The Filter attachment points used by the  
9 toolbar attachment manager and the menu attachment manager would use the  
10 same Explode attachment point as their data sources but would expose different  
11 collections of nodes, because they were filtering based on different criteria.

12       Consider Fig. 20 which incorporates an EDFHub attachment point 2000.  
13 This attachment point receives all of the EDFs and, as discussed above, combines  
14 them into a single XML list. The EDFHub then exposes the root nodes of all of  
15 the EDFs. The explode attachment point 2002 then exposes an XML list that  
16 contains all of the top level nodes for all of the EDFs. As an example, there may  
17 be multiple EDFs that each contain top level menu nodes, toolbar nodes,  
18 accelerator nodes and the like. Explode attachment point 2002 exposes an XML  
19 list that contains all of these top level nodes for all of the EDFs. Filter attachment  
20 point 2004 can then filter the XML list exposed by the explode attachment point  
21 2002 in accordance with any suitable parameters (i.e. filter on menu nodes, tool  
22 bar nodes, accelerator nodes and the like). The final explode attachment point  
23 2006 then exposes a list of the individual children nodes of the list exposed by the  
24 filter attachment point 2004. This list describes all of the specific features (of the  
25 particular type that were filtered) that have been added by all of the EDFs.

The table below lists a number of different attachment points that can be utilized in accordance with this described embodiment but many more can easily be created.

Attachment Point	Purpose
URL	Loads the URL and exposes the top level XML node as a member of the collection
Context	For every member, it gets the "expression" attribute and binds to it. If the expression evaluates to true, then the member is exposed.
EDF	Same as the URL AP, but also exposes a fabricated member with data to create an APP based on the URL and URN (which exists in the XML DOM).
Merge	Takes zero or more Attachment Points (of any type) and merges them together. The order and continuity of the original collections will be maintained.
Filter	Monitors a single Attachment Point and only exposes those nodes that match the specified name. The order of the original collection will be maintained.
Duplicate	Monitors a single Attachment Point and filters out any duplicates. A duplicate is defined to be having the same URN attribute. If no URN attribute is present then the node is exposed. Order of the original collection will be maintained.
Explode	Monitors a single Attachment Point and for every member exposes the children of that member as its members. The order of the original collection will be maintained as well as the order of the children within the nodes.
Link	Monitors a single Attachment Point and for every member looks for a URL attribute and creates a URL AP and

	merges it into itself. If the optional include Content is set to true, it will merge the original collection in as well.
Order	Monitors a single Attachment Point. For every member, it gets three attributes: id, before and after. Based on this information, it reorders the members as specified. If no ordering information is supplied, the order of the original collection will be maintained.
EDFHub	This Attachment Point is the central merge point that represents all the EDF points.

### **Conclusion**

The embodiments described above provide a platform solution that provides for customization and extensibility through a consistent and logical extensibility mechanism and object model that can be easily understood by third party developers. Internet-based downloads can be accomplished without a great deal of user intervention and without manipulating any user persisted settings. Extensions can be provided to a software platform or application program dynamically based upon the user's computing context.

Although the invention has been described in language specific to structural features and/or methodological steps, it is to be understood that the invention defined in the appended claims is not necessarily limited to the specific features or steps described. Rather, the specific features and steps are disclosed as preferred forms of implementing the claimed invention.

007290-AT66560

1    **CLAIMS**

2        1.    A software architecture comprising:  
3            multiple attachment points collectively arranged to filter data associated  
4            with files that describe software extensions; and  
5            multiple extension managers associated with the multiple attachment points  
6            and with respective feature types that can be added to a software platform by  
7            software extensions, the extension managers being configured to receive data from  
8            the multiple attachment points that pertains only to the feature type with which the  
9            extension manager is associated.

10  
11        2.    The software architecture of claim 1, wherein the attachment points  
12            are defined as predicate chains.

13  
14        3.    The software architecture of claim 1, wherein the attachment points  
15            filter XML data.

16  
17        4.    The software architecture of claim 3, wherein each feature type is  
18            associated with an XML tag.

19  
20        5.    The software architecture of claim 3, wherein each feature type is  
21            associated with an XML tag, at least some of the feature types comprising user-  
22            defined feature types.

1           6.     The software architecture of claim 1, wherein each attachment point  
2 exposes collections of ordered nodes.

3  
4           7.     The software architecture of claim 1, wherein each attachment point  
5 exposes collections of ordered XML nodes.

6  
7           8.     A computer embodying the software architecture of claim 1.

8  
9           9.     A software architecture comprising:  
10 a hub structure configured to:  
11                receive multiple different files that describe software extensions that  
12 can be added to a software platform;  
13                combine the multiple different files into a single exposable list; and  
14                expose the single exposable list to a filter structure that is configured  
15 to filter the list.

16  
17          10.    The software architecture of claim 9, wherein the hub structure  
18 receives multiple different XML files and exposes a list of XML nodes.

19  
20          11.    A computer embodying the software architecture of claim 9.

1        **12.**    A software architecture comprising multiple different attachment  
2 points each of which is configured to:

3            receive XML data that pertains to one or more software extensions that can  
4 be added to a software platform;

5            process the XML data to provide a list of XML nodes; and

6            expose the list of XML nodes.

7  
8        **13.**    The software architecture of claim 12, wherein the list of XML  
9 nodes is exposed to another attachment point.

10  
11       **14.**    The software architecture of claim 12, wherein the list of XML  
12 nodes can pertain to multiple different feature types that can be added by the one  
13 or more software extensions.

14  
15       **15.**    The software architecture of claim 12, wherein the list of XML  
16 nodes can pertain to multiple different features of particular feature types that can  
17 be added by the one or more software extensions.

18  
19       **16.**    The software architecture of claim 12, wherein the list of XML  
20 nodes can pertain to one or more of:

21            multiple different feature types that can be added by the one or more  
22 software extensions; and

23            multiple different features of particular feature types that can be added by  
24 the one or more software extensions.

1       **17.**    A computer embodying the software architecture of claim 12.

2  
3       **18.**    A software architecture comprising:

4       a hub structure configured to:

5               receive multiple different files that describe software extensions that  
6       can be added to a software platform;

7               combine the multiple different files into a single exposable list; and  
8               expose the single exposable list to a filter structure that is configured  
9       to filter the list;

10              a filter structure comprising multiple attachment points collectively  
11       arranged to filter data associated with the list exposed by the hub structure; and

12              multiple extension managers associated with the multiple attachment points  
13       and with respective feature types that can be added to a software platform by  
14       software extensions, the extension managers being configured to receive data from  
15       the multiple attachment points that pertains only to the feature type with which the  
16       extension manager is associated.

17  
18       **19.**    The software architecture of claim 18, wherein the hub structure  
19       receives multiple different XML files and exposes a list of XML nodes.

20  
21       **20.**    The software architecture of claim 19, wherein the list contains root  
22       node tags for all of the XML files.



1           **21.**   The software architecture of claim 19, wherein the XML files  
2 logically describe where a particular extension fits on the software platform.

3  
4           **22.**   The software architecture of claim 19, wherein the attachment points  
5 are defined as predicate chains.

6  
7           **23.**   The software architecture of claim 19, wherein an extension  
8 manager is notified whenever a extension comprising a feature type with which it  
9 is associated is added or removed from the software platform.

10  
11           **24.**   The software architecture of claim 19, wherein each feature type is  
12 associated with a particular XML tag.

13  
14           **25.**   A computer embodying the software architecture of claim 18.

15  
16           **26.**   A method of providing a software extension comprising:  
17 exposing an XML list that contains one or more nodes;  
18 processing the XML list to identify specific nodes that correspond to  
19 various feature types that can be added to a software platform; and  
20 notifying an extension manager that is associated with at least one feature  
21 type if a node that corresponds to that feature type is identified in the XML list.

1           **27.**    The method of claim 26, wherein said processing is accomplished  
2 by filtering the XML list using multiple attachment points that are defined as  
3 predicate chains.

4  
5           **28.**    The method of claim 27, wherein the individual attachment points  
6 receive XML data as an input and expose a list of XML nodes.

7  
8           **29.**    The method of claim 26, wherein said processing is accomplished  
9 by filtering on specific nodes.

10  
11          **30.**    The method of claim 26, wherein said processing is accomplished  
12 by exploding various nodes.

13  
14          **31.**    The method of claim 26, wherein said processing is accomplished  
15 by filtering on specific nodes and exploding various nodes.

16  
17          **32.**    One or more computer-readable media having computer-readable  
18 instructions thereon which, when executed by a computer, cause the computer to  
19 implement the method of claim 26.

20  
21          **33.**    A method of providing a software extension comprising:  
22            receiving XML data that pertains to a software extension that is to be added  
23 to a software platform;  
24            processing the XML data to identify XML nodes; and  
25

1 exposing an XML list that contains one or more nodes that are identified by  
2 said processing.

3  
4 **34.** The method of claim 33, wherein said receiving comprises receiving  
5 multiple XML files that pertain to different software extensions.

6  
7 **35.** The method of claim 34, wherein said processing comprises  
8 combining the multiple XML files into a single exposable list.

9  
10 **36.** The method of claim 33, wherein said processing comprises  
11 processing the XML data with one or more attachment points that are defined as  
12 predicate chains that filter the XML data.

13  
14 **37.** The method of claim 36, wherein at least one of the attachment  
15 points exposes a node.

16  
17 **38.** The method of claim 36, wherein at least one of the attachment  
18 points filters on a node.

19  
20 **39.** One or more computer-readable media having computer-readable  
21 instructions thereon which, when executed by a computer, cause the computer to  
22 implement the method of claim 33.

1           **40.**    A method of providing a software extension comprising:  
2           receiving multiple different files, each of which being associated with a  
3           different software extension and logically describing its associated software  
4           extension;  
5           combining the multiple different files in a single list;  
6           exposing portions of the list;  
7           processing exposed portions of the list to identify one or more feature types  
8           that are to be added to a software platform; and  
9           notifying an extension manager that is associated with a particular feature  
10          type.

11  
12          **41.**    The method of claim 40, wherein the multiple different files  
13          comprise XML files.

14  
15          **42.**    One or more computer-readable media having computer-readable  
16          instructions thereon which, when executed by a computer, cause the computer to  
17          implement the method of claim 40.  
18  
19  
20  
21  
22  
23  
24  
25

1 **ABSTRACT**

2       Methods and systems for network-based or Internet-based software delivery  
3 are described. In one embodiment, an application program or software platform  
4 resides on a client. The program or platform is configured so that it is extensible  
5 based on software extensions that are deliverable over a network such as the  
6 Internet. Various extensions can be developed by third party developers for  
7 incorporation into the program or platform. In the described embodiment,  
8 extension files that comprise a software extension are hosted on an Internet server.  
9 Additional files are also hosted on the Internet server and describe aspects of the  
10 software extension. These additional files include files that describe a logical  
11 attachment of the extension to the program or platform, as well as files that  
12 describe the location of the extension files. Extensions are incorporated on a client  
13 by navigating to a particular Internet site through which the extensions can be  
14 accessed. The files describing the extension files are downloaded on the client.  
15 These files tell the client where the particular extension is to be plugged in, as well  
16 as where to find the appropriate extension files and how to download them. The  
17 extension files are then downloaded and incorporated into the program or  
18 platform.

19  
20  
21  
22  
23  
24  
25

001230-286650

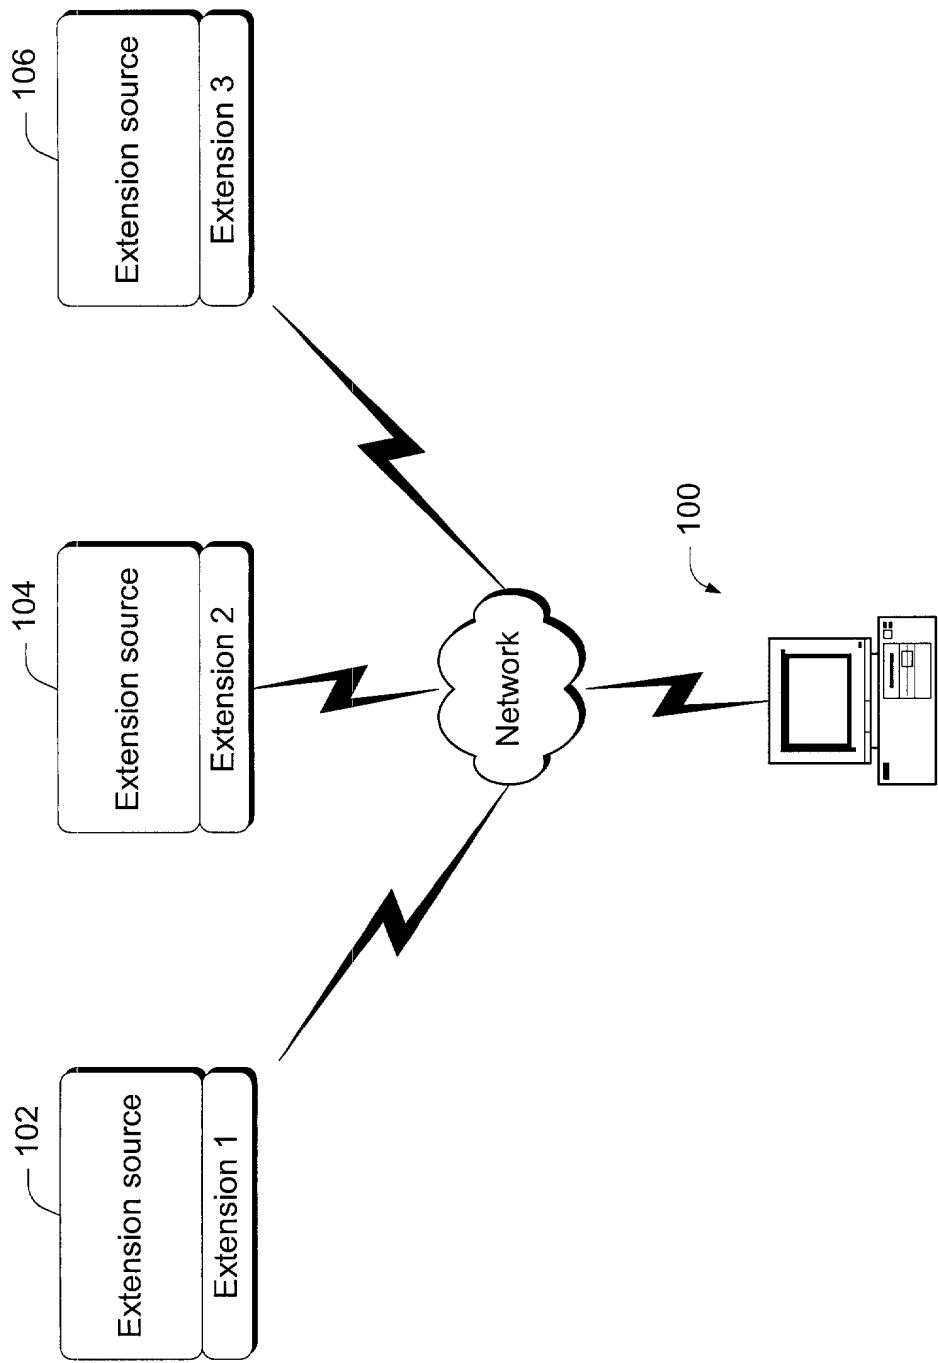



Fig. 1

Fig. 2



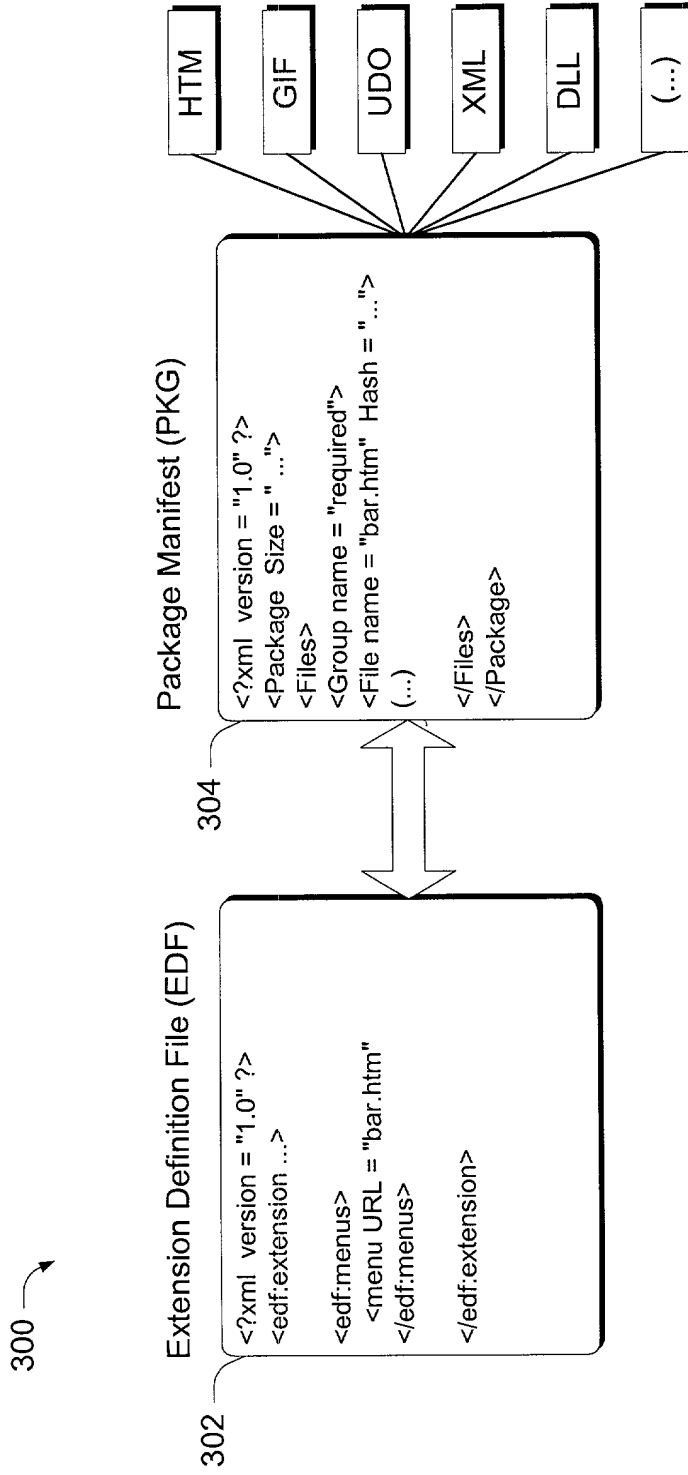


Fig. 3



Extension Tag

```
<edf:extension
urn = "urn:HelloNetDocs.netdocs.com"
name = "HelloNetDocs"
version = "1.0.0.0"
lastUpdate = "2000-01-15 8:43:52"
description = "Hello NetDocs"
>
(...)
</edf:extension>
```

Fig. 4

EDF Schema

```
<?xml...>
<!--info about the extensions and linked extensions -->
<edf:extension ...>

    <edf:edfLink ... />

<!--UI extensions -->
    <edf:toolbars> ... </edf:toolbars>
    <edf:accelerators> ... </edf:accelerators>
    <edf:menuItems> ... </edf:menuItems>
    <edf:commandGroups> ... </edf:commandGroups>
    <edf:themes> ... </edf:themes>

<!-- Objects, Components and Behavior Extensions -->
    <edf:Services> ... </edf:Services>

<!-- LocalStore Extensions -->
    <edf:contentClasses> ... </edf:contentClasses>
    <edf:subscriptions> ... </edf:subscriptions>

</edf:extension>
```

Fig. 5

00T290" 2T866560

600 →

Package Manifest

```
<?xml version "1.0" ?>
<Package Size= " ..." >
  <Files>
    <Group Name = "required".
      <File Name="bar.dll" />
      <File Name="foo.dll" Hash=" ..." />
    <Dependencies>
      <File Name="bar.dll" />
    </Dependencies>
    <COMClass ClassID=" ..." />
  </File>
  <File Name="myextension.edf" Hash="..." />
  <File Name="errorhandling.gif" Hash="..." />
</Group>
<Group Name = "offline">
  <File Name="bar.xml" Hash="..." />
  <File Name="abc.gif" Hash="..." />
  <File Name="xyz.gif" Hash="..." />
  <File Name="foo.htm" Hash="..." />
</Group>
<Group Name = "ondemand">
  <File Name="foo2.htm" Hash="..." />
</Group>
</Files>
</Package>
```

Fig. 6

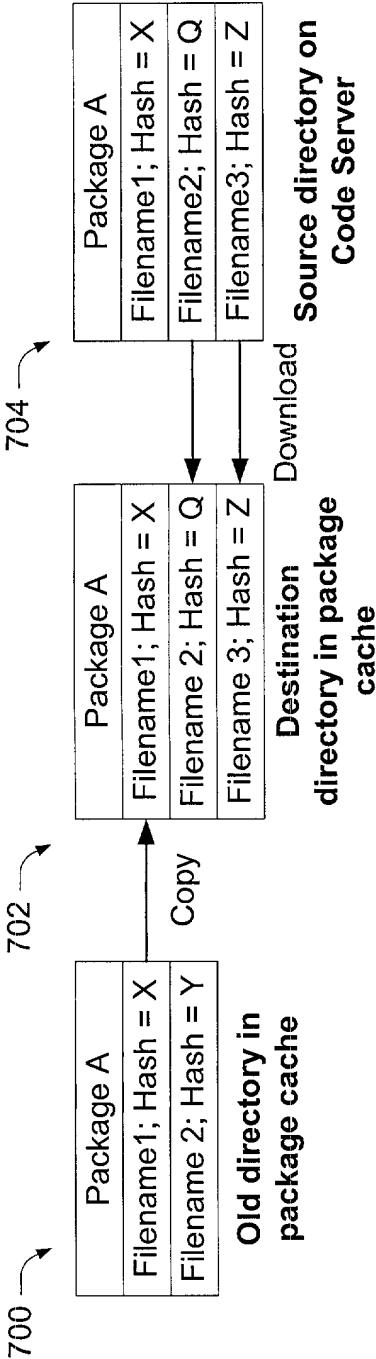


Fig. 7


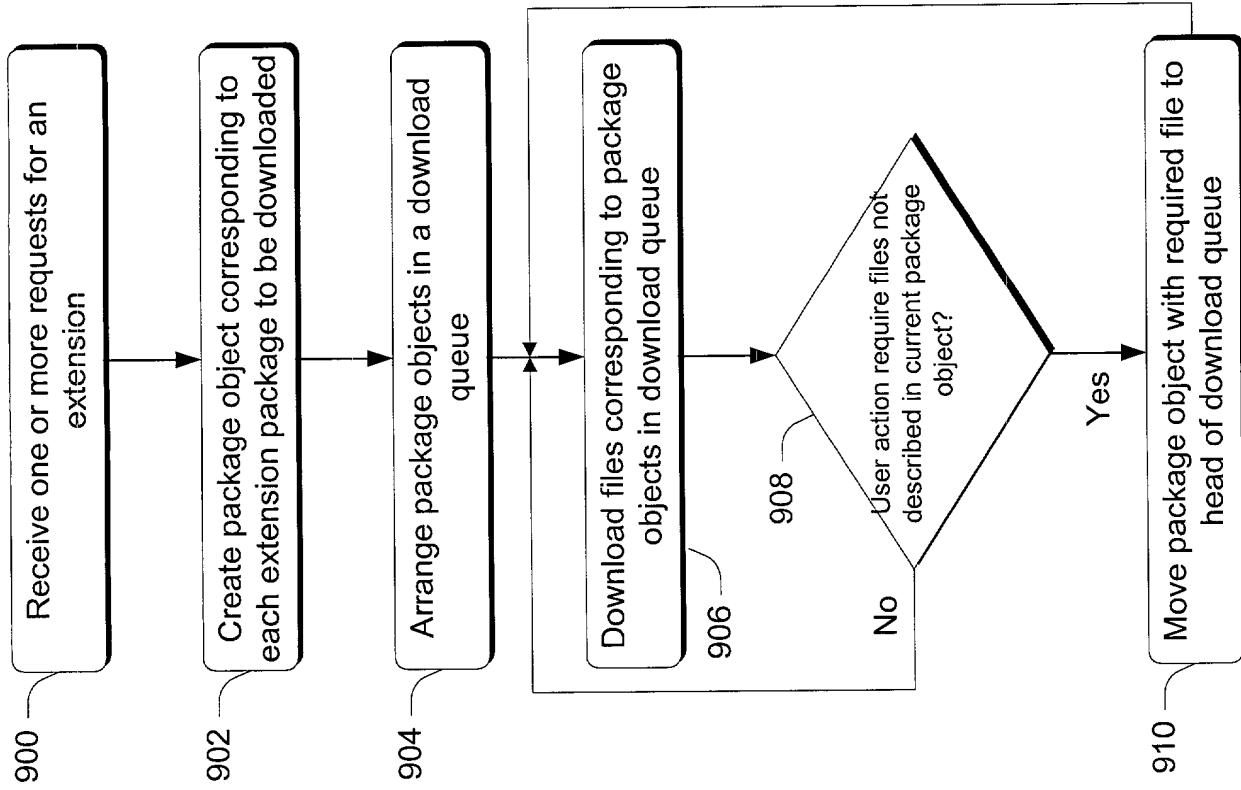
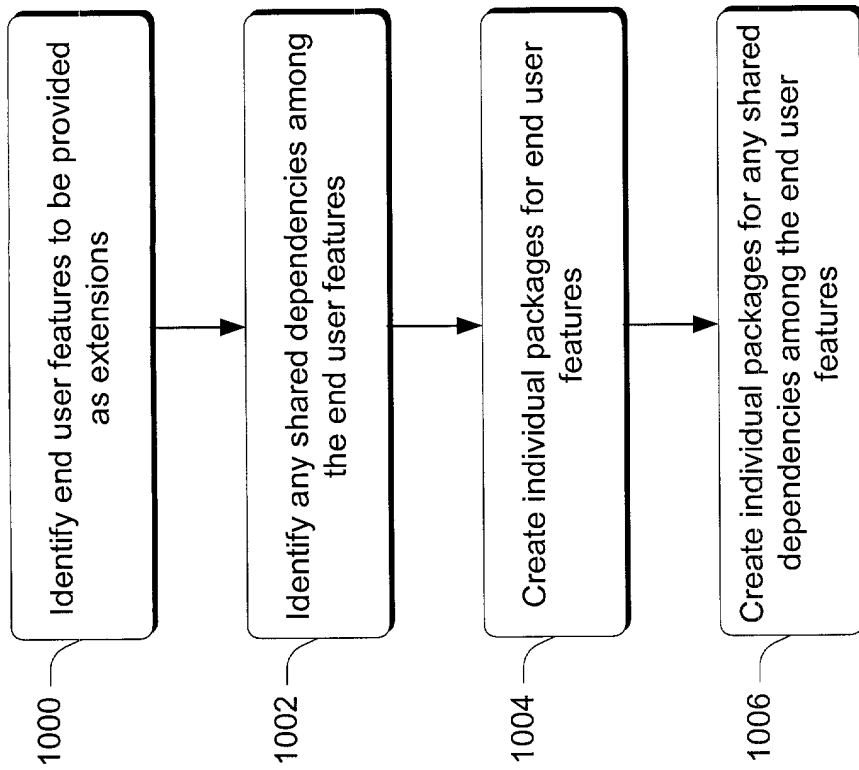



Fig. 8

Fig. 9



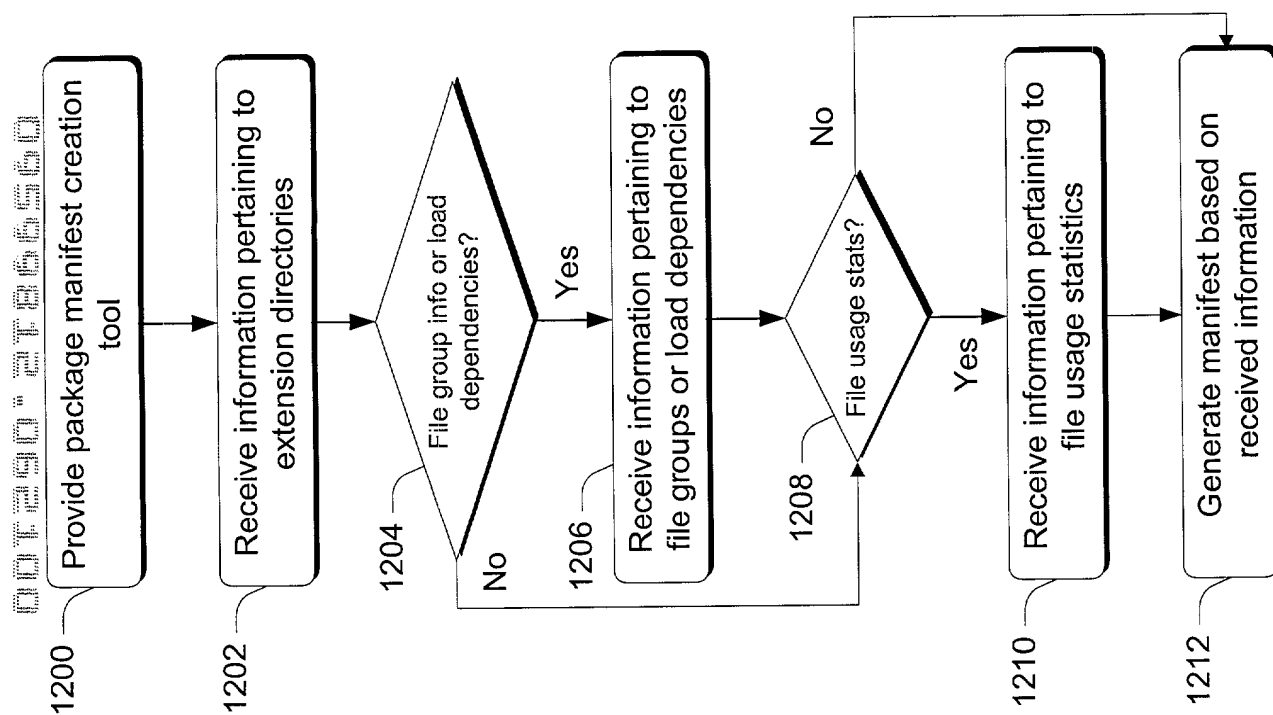



**Fig. 10**



*Fig. 11*

Fig. 12





*Fig. 13*



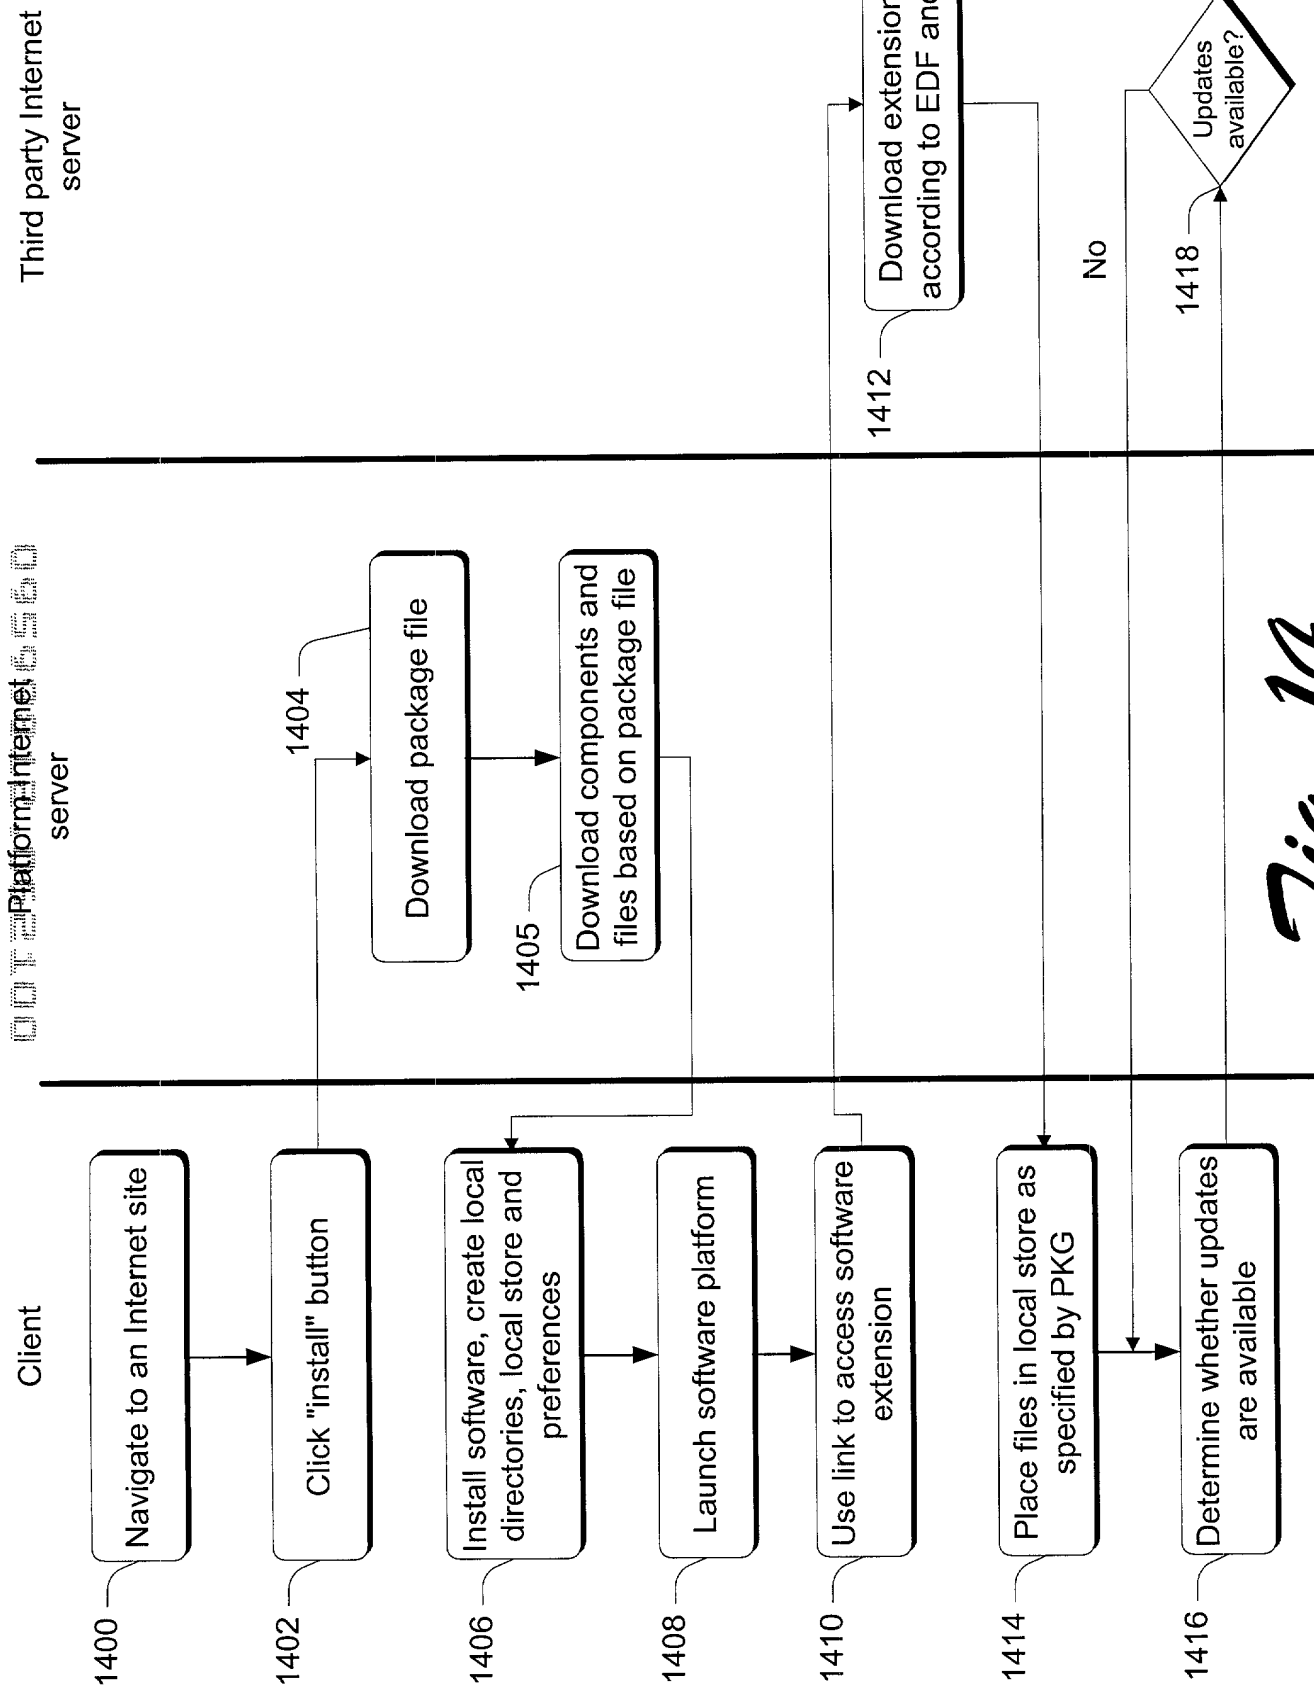


Fig. 14

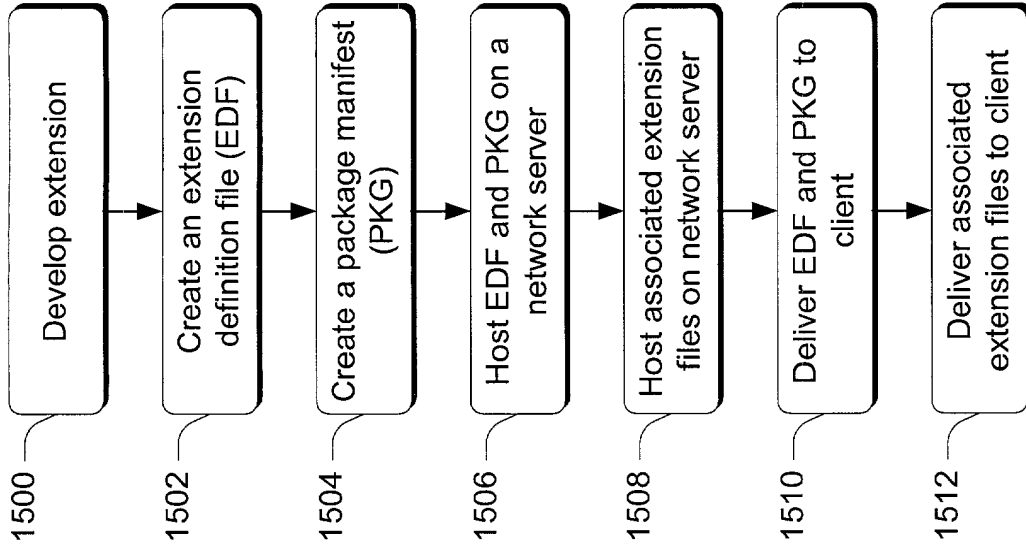


Fig. 15

```

<cat:extensionCatalog xmlns:cat="urn:schemas-microsoft-com:netdocs_edfcatalog.v1"
  lastUpdate="2000-01-15 8:43:52"
  description = "Catalog of extensions from hosted service"
>
<cat:extension extensionURN="urn:planner.netdocs.com"
  name="Planner"
  packageURN = "urn:planner.netdocs.com.ENv1"
  url="http://netdocs/extensions/planner/v1/EN/sc.cab"
  language="English"
  version="1"
  defaultLanguage="yes"
/>
<cat:extension extensionURN="urn:planner.netdocs.com"
  name="Planner"
  packageURN = "urn:planner.netdocs.com.FRv1"
  url="http://netdocs/extensions/planner/v1/FR/sc.cab"
  language="French"
  defaultversion="yes"
  version="1"
/>
<cat:extension extensionURN="urn:planner.netdocs.com"
  name="Planner"
  packageURN = "urn:planner.netdocs.com.ENv1.1"
  url="http://netdocs/extensions/planner/v1.1/EN/sc.cab"
  language="English"
  version="1.1"
  defaultLanguage="yes"
  defaultversion="yes"
/>
</cat:extensionCatalog>
    
```

Fig. 16





Fig. 17



*Fig. 18*

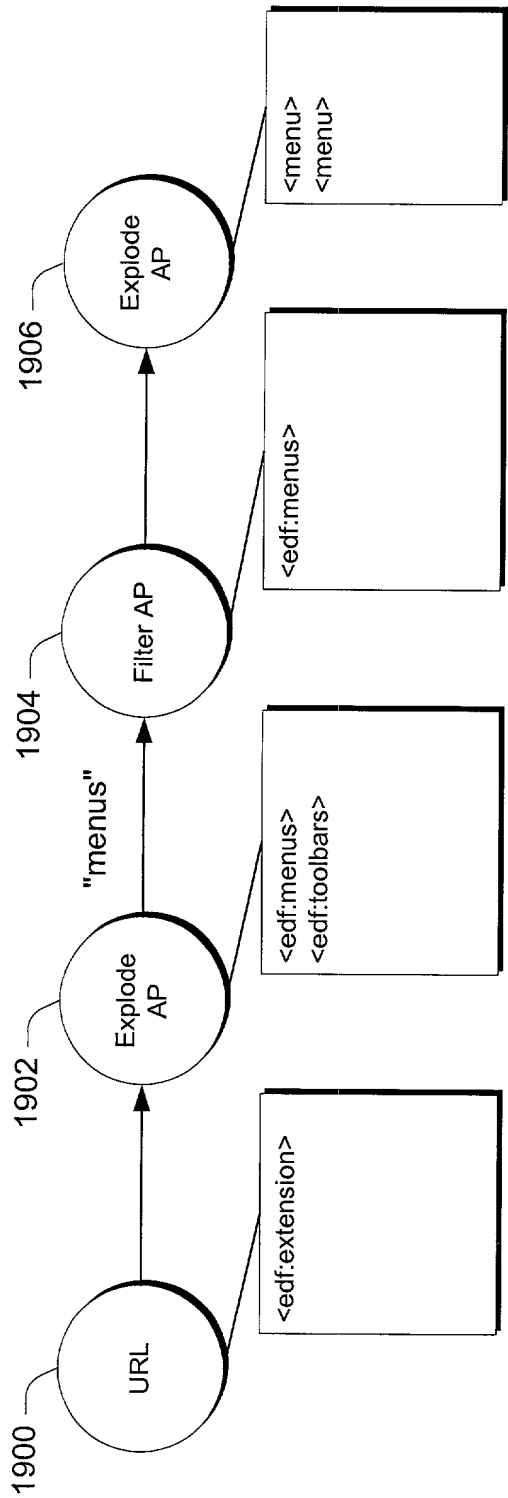


Fig. 19




Fig. 20